Realizable Deployment of Limited-Knowledge Robotic Inspectors for Nuclear Verification

Eric Lepowsky and David Snyder[1]

Abstract— **This work underscores the prospect and fundamental challenges of deploying robotic radiation detectors in nuclear safeguards and arms control applications, where safety and security concerns are of utmost importance. We propose a performant, minimal-knowledge approach which addresses information security constraints, with specific emphasis on the challenge of confirming the absence of undeclared radioactive sources. We demonstrate a random walk process which requires no *a priori* knowledge of the environment and certifies the absence or presence of sources without revealing any *a posteriori* information. We improve the convergence of this method by incorporating directional radiation measurements. The constraints of finite time and physical safety may be addressed by limited, rather than minimal, knowledge approaches, which illuminates a spectrum of limited-information contexts not yet thoroughly explored in this application space.**

I. INTRODUCTION

Nuclear safeguards and arms control are cornerstones of the broader global security mission. Safeguards, which are measures to verify that nuclear facilities are not misused and nuclear material is not diverted from peaceful uses, require verification throughout the nuclear fuel cycle [1], [2]. Arms control, which is concerned with limiting arms competition and regulating arsenals, is made possible through declarations, data exchange, and inspections to verify compliance with agreed upon limits [3]–[5]. Future agreements will likely require new verification approaches that minimize the need for access to sites, such as storage or dismantlement facilities, and treaty accountable items [6].

The introduction of robotics to the field of nuclear safeguards and treaty verification has the potential to be a paradigm shift in efficiency, effectiveness, and capability [7]–[10]. A "robotic inspector" enables remote inspections, wherein a human need not be present at the inspected site, thereby reducing safety and security concerns. Inspection tasks may include a variety of typically human-based measures, including verification of labels and seals, counting of objects, and, of principle interest here, radiation measurements. The eventual transition from human-in-the-loop to autonomous inspections will require a high level of confidence in the systems and processes [11]. Although robotic systems have been proposed, and in some cases deployed, for applications in nuclear disaster response and facility monitoring [12]–[14], there remain significant hurdles for deployment in nuclear safeguards and treaty verification. Here we highlight relevant considerations for implementing a robotic detector, particularly considering the design of algorithms subject to regulation and certification requirements.

[1]Mechanical and Aerospace Engineering, Princeton University, NJ, USA. E. L. and D. S. contributed equally.
Contact: {lepowsky,dasnyder}@princeton.edu

Certification and Authentication. Regulation and certification is a fundamental concept in this application space. Radiation equipment must be certified and authenticated [15]–[17]. Certification verifies that the system does what it's designed to do (and nothing more). Authentication requires verifying that the system provides genuine data.

Physical Safety. Safety verification is also crucial when considering the operation of remote or autonomous systems in environments containing sensitive objects. Potential high-risk environments encountered in safeguards and arms control include navigating rows of gas centrifuges for uranium enrichment and searching storage facilities with weapon components or delivery vehicles. In both scenarios, collisions could be catastrophic for health and safety, potential damage to infrastructure, and diminished trust from the host party.

Privacy. A suitable solution for deploying robotic inspectors would require that no sensitive information – which may include images, dimensions, radiation measurements, etc. – is unnecessarily or inadvertently revealed. Reported methods in robotic radiation detection primarily perform source localization or mapping; as such, many of these techniques assume either full *a priori* knowledge of the search environment or, by consequence of the algorithmic design, reveal full knowledge of the environment configuration or its radiation field *a posteriori* [18]–[22]. We also note that current approaches do not consider the inverse problem of confirming the *absence* of sources. Toward resolving the aforementioned security concerns, we focus on absence confirmation [23], in addition to source localization.

II. PROBLEM DEFINITION

Assume that a physical, bounded environment is declared to contain no radioactive sources. The associated verification task is to actively explore the two-dimensional obstacle-filled space such that the robot certifies the absence of sources (or presence, in the case of non-compliance). Ideal certification methods should allow for provable correctness (the robot will return the correct decision) and/or provable privacy (the robot's capacity to "leak" information is minimal).

The algorithm we propose takes inspiration from randomized, sampling-based motion planners [24], [25] and out-of-distribution detection [26], with an emphasis on scalar task-relevant detection as in [27]. When measurements are consistent with source absence, the robot moves according to a ("reference") random walk that explores the space; if the reading is consistent with source presence, it moves according to a different ("out-of-distribution") random walk. Detection of a shift is accomplished by Kolmogorov-Smirnov (KS) testing [28] of the composite (realized) action distribution. Further, because the actions depend only on the detected counts, the resulting distribution over actions for any source-free map is theoretically identical; such a property can help provide evidence that the algorithm *never* constructs a usable representation of the space during operation.

III. Minimal Knowledge Baseline

Practically, we assume that a robot is equipped with the capacity to rotate and translate in a controllable manner (e.g., via encoders on a wheeled system), to detect imminent collisions in a non-destructive fashion, and to accurately acquire radiation measurements. Such a system can run Alg. 1 (see App. V-A), wherein the robot "slows down" (i.e., takes smaller step sizes) when anomalously high counts are detected, such that a different distribution over step sizes is achieved. The proposed approach derives the performance properties of the following lemma:

Lemma 1 (Minimal Knowledge Properties): Assume that Alg. 1 is run on a traversable map E of outer dimensions (L_x, L_y). In the discrete case, if E can be observably discretized to tolerance $\epsilon > 0$, then the expected coverage time, for any initial configuration, is bounded by $\tilde{\mathcal{O}}(\frac{1}{\epsilon^4})$. Further, the false positive rate is no greater than p^*.

Proof: The claim follows by initially observing that Alg. 1 is a linear combination of uncorrelated random walks on a graph $\mathcal{G}_E(V, \mathcal{E})$ in two dimensions. Results from [29], [30] give an upper bound on expected coverage time of $2|V||\mathcal{E}|$. Because the connecting edges of a random walk are local, we can bound the number of edges by $|\mathcal{E}| \leq k|V|$, thereby giving a coverage bound of $2k|V|^2$. Noting that the discretization has $L_x L_y/\epsilon^2$ nodes completes the result. The calibration of the algorithm to p^* relies on controlling for the pre-selected number of tests n via the testing threshold of p^*/n. The resulting union bound is robust to correlation in the input data to the sequential KS tests. ■

The general extension of covering time results for two-dimensional random walks from discrete to continuous space is given by [31], meaning that the above result can generalize to models in the continuous setting. Verifying this property and the privacy characteristics of the minimal-information baseline is ongoing work.

IV. Increasing Information

The addition of an omnidirectional radiation detector, such as [32], [33], enables a similar pseudo-random walk policy. In the absence of a prevailing source, the direction is random due to the Poisson-distributed counting statistics; if a source is present, the robot more expeditiously switches to the "slowed down" policy. The inclusion of directionality manifests as noisy gradient ascent, with similar functionality as infotaxis [34], but with a less detailed sensor model. Fig 1 demonstrates the convergence of the proposed approach with and without the addition of directional information. Sample random walk trajectories are visualized in Fig. 3, as are the cumulative density functions (Fig. 4) corresponding to the same observed environments (see App. V-C).

As another possible extension, distribution testing may also be used to confirm that an environment remains unchanged by redefining the reference action distribution. This introduces the concept of template matching, wherein "absence" is reframed in terms of the absence of deviations, rather than the absence of absolute sources.

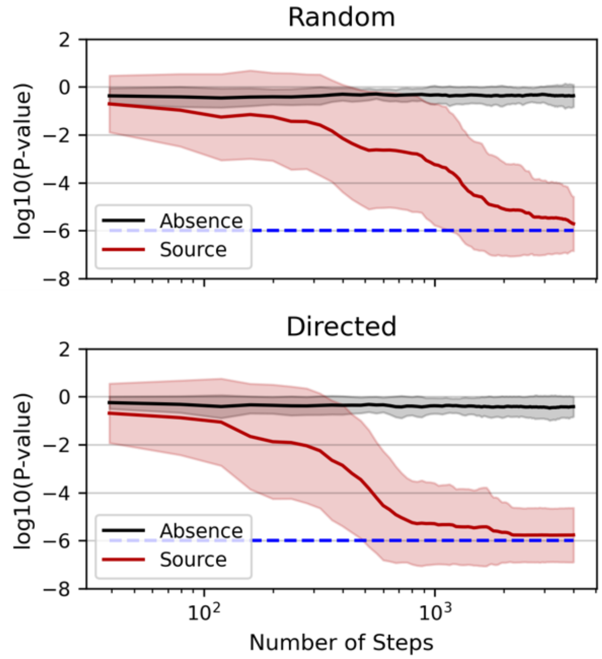

Fig. 1. Evolution of floored KS test p-values for baseline (top) and directed (bottom) algorithms, averaged over 50 seeds, for source presence (red) and absence (black). In all settings, a log-significance below -6 served as the cutoff for detection. For both cases, source-absence is correctly identified. While both variants show steady convergence for the source-presence case, the directed algorithm demonstrates faster convergence due to its utilization of some additional information in the radiation sensing model.

Further advancements may also be made by relaxing the constraints of no *a priori* or *a posteriori* knowledge of the search environment. While theoretically sound, a random walk process may still be impractical for time-efficient inspections. There is a benefit to including limited contextual knowledge to expedite the inspection process. Here, we offer just one example of a more "knowledgeable" approach. Particle filtering, commonly used for radioactive source localization [35]–[38], can be adapted to absence confirmation by extending the range of admissible source intensities to zero. Sufficient coverage of the accessible environment is necessary to eliminate weak-source hypotheses with high confidence; unfortunately, efficient coverage algorithms typically require knowledge of the environment [39]–[41]. To improve the convergence efficiency while also minimizing requisite environment knowledge, LiDAR may be used to effectively eliminate nonzero-intensity particles in free space, under the assumption that a radioactive source must be bound to a physical feature and non-compliant objects in the applications of interest would be sufficiently substantial. This should enable more efficient exploration than a random walk by moving toward where potential sources may be located. LiDAR also enables simple collision avoidance, even without full map generation, providing a degree of safety which would otherwise be impossible without alternative proximity sensors or odometry.

V. APPENDIX

A. Minimal Knowledge Exposition

The minimal-knowledge random walk algorithm (without directional sensing) is presented below in the paradigm of continuous space. It can be directly modified to the discrete case by discretizing acceptable orientations (e.g., 4- or 8-connecting graphs) and step sizes.

Algorithm 1 Random walk absence confirmation.

Input: Background count rate B, outer dimensions L_x, L_y, confidence parameter p^*, run time T, test count n, threshold level z, constants $0 \leq c_L < c_U$, reference V_r
Initialize $\underline{p} = 1.0$, $x_0, y_0, \theta_0 \sim \mathcal{U}(L_x, L_y, 2\pi)$, $t = 1$, $V_e = \{\emptyset\}$
while $t \leq T$ **do**
  $N_t \sim h(x_t, y_t; E)$ {Measurement}
  $c \leftarrow c_L + (c_U - c_L)\mathbb{1}[N_t \leq B + z\sqrt{B}]$

  $ds, d\theta \sim \mathcal{U}(c, 2\pi)$ {Step Length, Rotation}
  Rotate by $d\theta$ rad. and move forward ds distance
  Append ds to memory V_e
  if $t \equiv 0 \pmod{T/n}$ **then**
    $\underline{p} = \min\{\underline{p}, \mathbf{KS}(V_e, V_r)\}$
  end if
  if $\underline{p} \leq p^*/n$ **then**
    return 1 {Anomaly detected}
  end if
end while
return 0 {Absence confirmed}

B. Additional Random Walk Variant

As an additional point of comparison, Fig. 2 reproduces the results of Fig. 1 for a random walk scheme inspired by Lévy flight [42]. After every 100 time steps, the robot "teleports" to a uniformly random sampled position within the bounding box. While this does not follow the Lévy distribution, it effectively emulates the multi-scale step lengths that are characteristic of Lévy flights.

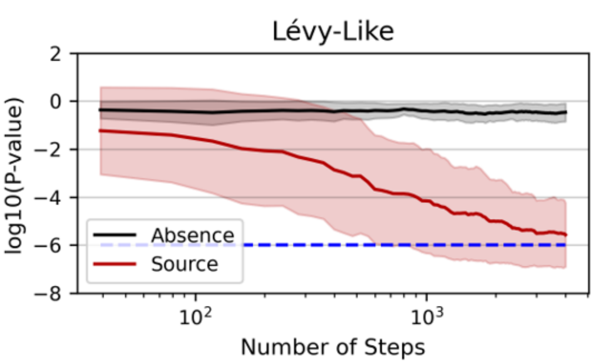

Fig. 2. Evolution of floored KS test p-values for Lévy-inspired algorithm, averaged over 50 seeds, for source presence (red) and absence (black). A log-significance below -6 served as the cutoff for detection. The Lévy-inspired algorithm demonstrates convergence speed between that of the algorithms shown in Fig. 1.

C. Preliminary Random Walk Demonstration

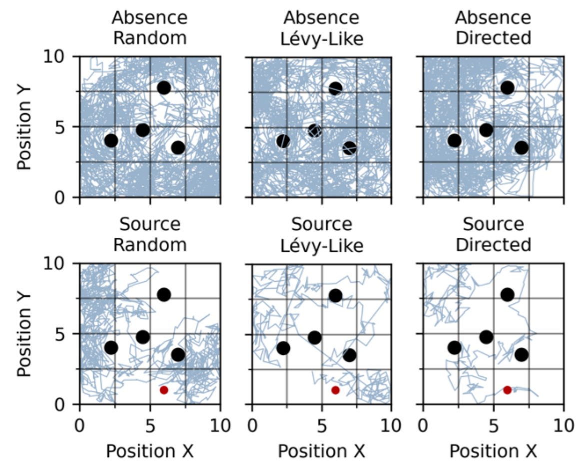

Fig. 3. Representative trajectories of the minimal-knowledge robotic radiation detector in obstacle-filled environments with (below) and without (above) a source. For each case, three variants of the proposed random walk algorithm are applied: the robot moves in a uniformly random direction for each step (left); the robot follows a Lévy-inspired sampling scheme, allowing it to infrequently "teleport" long distances (middle); the robot moves in the apparent direction provided by radiation sensor measurements (right). All algorithms demonstrate qualitatively broad covering behavior in their respective absence cases, but the directed case more expeditiously finds the source. Note that the trajectories shown would not be feasible to reconstruct under the proposed algorithm, since only the step size is stored; without storing the measured direction or position, it is impossible to construct a meaningful map representation.

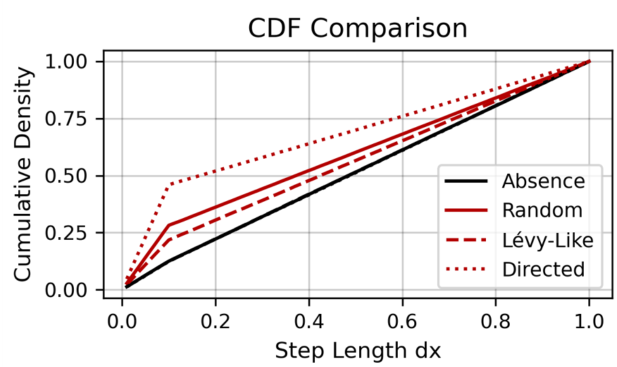

Fig. 4. Cumulative density function over step lengths, averaged over 50 seeds, for source presence (red) and absence (black), corresponding to the environments tested in Fig. 1 and Fig. 2. Note that the black dashed and dotted lines are nearly identical to the solid black line, which is indicative of the equivalency of the random walk methods in a source-absence environment. The directed algorithm demonstrates faster convergence due to its utilization of some additional information in the sensing model, which causes a larger shift in the realized action distribution.

ACKNOWLEDGEMENTS

We thank Alexander Glaser, Robert J. Goldston, and Anirudha Majumdar for their feedback and support. This work has been supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2039656.

REFERENCES

[1] *IAEA Safeguards: Serving Non-Proliferation*, International Atomic Energy Agency, 2018.

[2] J. Carlson, V. Kuchinov, and T. Shea, *The IAEA's Safeguards System as the Non-Proliferation Treaty's Verification Mechanism*, May 2020.

[3] J. Fuller, "Verification on the Road to Zero: Issues for Nuclear Warhead Dismantlement," Arms Control Today, December 2010.

[4] C. Comley, M. Comley, P. Eggins, G. George, S. Holloway, M. Ley, P. Thompson, and K. Warburton, *Confidence, Security & Verification, The Challenge of Global Nuclear Weapons Arms Control*, AWE/TR/2000/001, Atomic Weapons Establishment, Aldermaston, United Kingdom, 2000.

[5] *Radiation Detection Equipment: An Arms Control Verification Tool*, Product No. 211P, Defense Threat Reduction Agency, Fort Belvoir, VA, October 2011.

[6] National Academies of Sciences, Engineering, and Medicine, *Nuclear Proliferation and Arms Control Monitoring, Detection, and Verification: A National Security Priority: Interim Report*, The National Academies Press, 2021.

[7] F. F. Dean, *ROBIN: A Way to Collect In-Plant Safeguards Data with Minimal Inspector Access*, SAND82-1588C, Sandia National Laboratories, 1982.

[8] K. Robertson, R. Stohr, A. Elfes, P. Flick, A. Sokolov, D. Finker, and C. Everton, "The IAEA Robotics Challenge – Demonstrating Robots for Safeguards Inspections," IAEA Symposium on International Safeguards: Building Future Safeguards Capabilities, IAEA-CN-267/215, 2018.

[9] F. E. Schneider and D. Wildermuth, "Real-World Robotic Competitions for Radiological and Nuclear Inspection Tasks," *20th International Carpathian Control Conference (ICCC)*, pp. 1-6, 2019.

[10] B. Bird, A. Griffiths, H. Martin, E. Codres, J. Jones, A. Stancu, B. Lennox, S. Watson, and X. Poteau, "A Robot to Monitor Nuclear Facilities: Using Autonomous Radiation-Monitoring Assistance to Reduce Risk and Cost," *IEEE Robotics & Automation Magazine*, vol. 26, no. 1, pp. 35-43, 2019.

[11] M. Fisher, R. C. Cardoso, E. C. Collins, C. Dadswell, L. A. Dennis, C. Dixon, M. Farrell, A. Ferrando, X. Huang, M. Jump, G. Kourtis, A. Lisitsa, M. Luckcuck, S. Luo, V. Page, F. Papacchini, and M. Webster, "An Overview of Verification and Validation Challenges for Inspection Robots," *Robotics*, vol. 10, no. 2, 67, 2021.

[12] R. Smith, E. Cucco, and C. Fairbairn, "Robotic Development for the Nuclear Environment: Challenges and Strategy," *Robotics*, vol. 9 no. 4, 94, 2020.

[13] M. Chiou, G. T. Epsimos, G. Nikolaou, P. Pappas, G. Petousakis, S. Mühl, and R. Stolkin, "Robot-Assisted Nuclear Disaster Response: Report and Insights from a Field Exercise," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.

[14] I. Tsitsimpelis, C. J. Taylor, B. Lennox, and M. J. Joyce, "A Review of Ground-Based Robotic Systems for the Characterization of Nuclear Environments," *Progress in Nuclear Energy*, vol. 111, pp. 109-124, 2019.

[15] J. Brotz, *Trusting Embedded Hardware and Software in Treaty Verification Systems*, SAND2019-5077C, Sandia National Laboratories, 2019.

[16] *Food-for-Thought Paper: Equipment Authentication for Nuclear Dismantlement Monitoring*, Working Group 3: Technical Challenges and Solutions, International Partnership for Nuclear Disarmament Verification, 2017.

[17] D. W. MacArthur, D. K. Hauck, and M. K. Smith, *Confirmation of Nuclear Treaty Limited Items: Pre-dismantlement vs. Post-dismantlement*, LA-UR-13-23004, Los Alamos National Laboratory, 2013.

[18] F. Gagliardi, "Integration of Independent NDA Techniques within a SLAM-based Robotic System for Improving Safeguards Standard Routines: A Review of the Current Status and Possible Future Developments," *ESARDA Bulletin*, vol. 64, no. 2, pp. 10-21, 2022.

[19] D. Hellfeld, T. H. Y. Joshi, M. S. Bandstra, R. J. Cooper, B. J. Quiterm, and K. Vetter, "Gamma-Ray Point-Source Localization and Sparse Image Reconstruction Using Poisson Likelihood," *IEEE Transactions on Nuclear Science*, vol. 66, no. 9, pp. 2088-2099, 2019.

[20] F. Mascarich, P. D. Petris, H. Nguyen, N. Khedekar, and K. Alexis, "Autonomous Distributed 3D Radiation Field Estimation for Nuclear Environment Characterization," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2163-2169, 2021.

[21] A. West, I. Tsitsimpelis, M. Licata, A. Jazbec, L. Snoj, M. J. Joyce, and B. Lennox, "Use of Gaussian process regression for radiation mapping of a nuclear reactor with a mobile robot," *Scientific Reports*, vol. 11, 13975, 2021.

[22] N. A. Abd Rahman, K. S. M. Sahari, N. A. Hamid, and Y. C. Hou, "A coverage path planning approach for autonomous radiation mapping with a mobile robot," *International Journal of Advanced Robotic Systems*, vol. 19, no. 4, 2022.

[23] E. Lepowsky, J. Jeon, and A. Glaser, "Confirming the Absence of Nuclear Warheads via Passive Gamma-Ray Measurements," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 990, 164983, 2021.

[24] S. M. LaValle and J. J. Kuffner, "Randomized Kinodynamic Planning," *International Journal of Robotics Research*, vol. 20, no. 5, pp. 378-400, 2001.

[25] J. J. Kuffner and S. M. LaValle, "RRT-Connect: An Efficient Approach to Single-Query Path Planning," *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2, pp. 995-1001, 2000.

[26] , M. Basseville, "Detecting Changes in Signals and Systems – A Survey," *Automatica*, vol. 24, no. 3, pp. 309-326, 1988.

[27] A. Farid, S. Veer, and A. Majumdar, "Task-Driven Out-of-Distribution Detection with Statistical Guarantees for Robot Learning," *Proceedings of the 5th Conference on Robot Learning (PMLR)*, vol. 164, pp. 970-980, 2022.

[28] A. N. Kolmogorov, "Sulla Determinazione Empirica di Una Legge di Distribuzione," *Giornale dell'Istituto Italiano degli Attuari*, vol. 4, pp. 83-91. 1933

[29] D. Aldous, "An Introduction to Covering Problems for Random Walks on Graphs," *Journal of Theoretical Probability*, vol. 2, pp. 87-99, 1989.

[30] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovasz, and C. Rackoff, "Random Walks, Universal Traversal Sequences, and the Complexity of Maze Problems," *20th Annual Symposium on Foundations of Computer Science (SFCS)*, pp. 218-223, 1979.

[31] A. Dembo, Y. Peres, J. Rosen, and O. Zeitouni, "Cover Times for Brownian Motion and Random Walks in Two Dimensions," *Annals of Mathematics*, vol. 160, pp. 433-464, 2004.

[32] F. Mascarich, C. Papachristos, T. Wilsonm, and K. Alexis, "Distributed Radiation Field Estimation and Informative Path Planning for Nuclear Environment Characterization," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2318-2324, 2019.

[33] E. Lepowsky, M. Kütt, S. Aslam, H. Fetsch, S. Snell, A. Glaser, and R. J. Goldston, "Experimental Demonstration and Modeling of a Robotic Neutron Detector with Spectral and Directional Sensitivity for Treaty Verification," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1041, 167362, 2022.

[34] M. Vergassola, E. Villermaux, and B. Shraiman, "'Infotaxis' as a strategy for searching without gradients," *Nature*, vol. 445, pp. 406-409, 2007.

[35] B. Ristic, M. Morelande, and A. Gunatilaka, "A Controlled Search for Radioactive Point Sources," *11th International Conference on Information Fusion*, pp. 1-5, 2008.

[36] J. Huo, M. Liu, K. A. Neusypin, H. Liu, M. Guo, and Y. Xiao, "Autonomous Search of Radioactive Sources through Mobile Robots," *Sensors*, vol. 20, no. 12, 3461, 2020.

[37] Y. Liu, Y. Xuan, D. Zhang, and S. Zou, "Localizing unknown radiation sources by unscented particle filtering based on divide-and-conquer sampling," *Journal of Nuclear Science and Technology*, vol. 59, no. 9, pp. 1149-1161, 2022.

[38] M. Ling, J. Huo, G. V. Moiseev, L. Hu, and Y. Xiao, "Multi-robot collaborative radioactive source search based on particle fusion and adaptive step size," *Annals of Nuclear Energy*, vol. 173, 109104, 2022.

[39] Y. Cao, Y. Han, J. Chen, X. Liu, Z. Zhang, and K. Zhang, "Optimal Coverage Path Planning Algorithm of the Tractor-formation Based on Probabilistic Roadmaps," *IEEE International Conference on Unmanned Systems and Artificial Intelligence (ICUSAI)*, pp. 27-32,2019.

[40] Z. Khanam, S. Saha, D. Ognibene, K. McDonald-Maier, and S. Ehsan, "An Offline-Online Strategy for Goal-Oriented Coverage Path Planning using A Priori Information," *14th IEEE International Conference on Industry Applications (INDUSCON)*, pp. 874-881, 2021.

[41] S. A. Sadat, J. Wawerla, and R. Vaughan, "Fractal trajectories for online non-uniform aerial coverage," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2971-2976, 2015.

[42] G. M. Viswanathan, V. Afanasyev, S. V. Buldyrev, S. Havlin, M. G. E. da Luz, E. P. Raposo, and H. E. Stanley, "Lévy Flights in Random Searches," *Physica A: Statistical Mechanics and its Applications*, vol. 282, no. 1–2, pp. 1-12, 2000.